REMARKS

Claims 1-4, 8-14, 18-20, and 22-31 are now pending in the application. In response to the Office Action mailed January 16, 2008, Applicant respectfully requests reconsideration. Claims 1-4, 8-14, 18-20, 22-31 were previously pending in this application. Claims 1, 8-9, 11, 19, and 26 have been amended. Claims 1-4, 8-14, 18-20, 22-31 are pending for examination with Claims 1, 11, and 19 being independent. The application is believed to be in condition for allowance.

Summary of Embodiments of Applicant's Invention

An example of one embodiment of Applicant's invention is described below to highlight some aspects of the invention. It should be appreciated that the description below is merely an example of one of many embodiments that fall within the scope of Applicant's claims and is provided merely for the purpose of highlighting some aspects of Applicant's invention.

The application relates to systems and methods for providing differentiated classes of storage to clients accessing a storage system. This is accomplished by determining a level of performance for the storage locations within the system and then partitioning locations into regions as determined by their different levels of performance. In some embodiments, for example, a performance process measures the performance of storage locations by making experimental read and write operations across the logical block name space, and uses the measurements to determine whether various locations can be aggregated into regions (*See,* e.g., Specification page 7, lines 4-11). A mapping process maps and aggregates the logical block names of locations having an identical level of performance (i.e., the partitioned locations) to a section of the logical block name space (i.e., thereby creating different storage pools providing different classes of storage). For example, in some embodiments the system assigns different RAID levels to different regions based on the determined levels of performance of the locations within the regions (*See,* e.g., Specification page 2, lines 17-19 and page 4, lines 9-11; and Figure 4). Clients accessing the system can utilize the storage pool with the appropriate performance level needed to carry out the desired class of service (*See,* e.g., Specification page 11, lines 5-8).

For example, one client may utilize a RAID 10 service, and another may utilize a RAID 5 service (*See,* e.g., Specification page 9, line 22 to page 10, line 2).

Specification Objections

The Examiner objected to the Specification at page 1, lines 7-11, as claiming priority to U.S. Provisional Application No. 60/441810. The Examiner states that the provisional application fails to provide adequate support or enablement in the manner provided by the first paragraph of 35 U.S.C. 112 for one or more claims of this application. Applicant respectfully disagrees.

In the provisional application, support can be found at least in the figures on pages 23/26, 24/26, 25/26 (showing Pools A, B, and C; and RAIDs 5, 10, and 50), 26/26, and 7/11 and the related text. For example, the text on pages 14-15 of the specification states that the clients 12 will have need of the resources partitioned across the server group 116. Accordingly, each of the clients 12 will send requests to the server group 116. The clients 12 typically act independently, and as such, the client load placed on the server group 116 will vary over time. In a typical operation, a client 12 will contact one of the servers, for example, server 161, to access a resource, such as a data block, page (comprising a plurality of blocks), file, database table, application, or other resource. The contacted server 161 itself may not hold or have control over the requested resource. However, in a preferred embodiment, the server group 116 is configured to make all the partitioned resources available to the client 12 regardless of the server that initially receives the request. For illustration, Figure 7 shows two resources, one resource 180 that is partitioned over all three servers (servers 161, 162, 163) and another resource 170 that is partitioned over two of the three servers. In the application of the system 110 being a block data storage system, each resource 170 and 180 may represent a partitioned block data volume.

Moreover, for example on page 16/26 of the Appendix (EQLC-PXX-005) (copy attached hereto as Exhibit A), the text describes how differentiated pools of storage offer different performance across their LBN space.

> In a block storage environment with multiple devices, differentiated pools of storage are created by applying multiple RAID techniques simultaneously on each device. These pools of storage differ in regards to performance and reliability. Adaptive load balancing can be used to place data in pools to optimize service for a

single or multiple consumers of storage services. It is stated there, for example:

a) Storage devices offer different performance across their LBN space: some LBN ranges have higher bandwidth for read and/or write, some LBN ranges have higher throughput for small and/or random operations, some LBN ranges have slower bandwidth, and some LBN ranges have slower throughput.

b) RAID levels have different performance characteristics for bandwidth and throughput for read and write operations, as well as in both normal mode and degraded operating modes.

c) Device performance differences across LBN space, and performance and reliability differences of RAID techniques, can be combined of to create storage pools that offer different classes of service across a common set of storage devices to the consumers of storage services.

Adaptive load balancing provides the ability to adjust data placement in pools for continuous optimal performance.

The technology to be protected is – minimally – :

1) The use of multiple raid levels concurrently on each device to create differentiated classes of storage service.

2) The use of multiple raid levels concurrently on each device in combination with device LBN performance differences to create differentiated classes of storage service.

3) Load balancing storage services across multiple differentiated classes of storage service:

    a. Automatic, adaptive load balancing of a single or multiple storage services (volumes or file systems) across pools via
        i. Volume based allocation
        ii. Extent based allocation
        iii. Page-based allocation

In addition, Figs. 5 and 6 of the Appendix, at pages 22-23/26 (also attached in Exhibit A) illustrate a performance process and pooled storage board or performance levels (A, B, C) having

performance differences across the LBN space. Accordingly, withdrawal of this objection is respectfully requested.

Rejection of Claims Under 35 U.S.C. § 102

Claims 1-4, 8-14, 18-20, and 22-31 were rejected under 35 U.S.C. § 102(e) as being anticipated by Dimitri et al., U.S. Patent No. 6,839,802 ("Dimitri").

In response, Applicant respectfully traverses the rejection by stating that Dimitri does not disclose or suggest all of the elements of the independent claims.

In brief, Dimitri describes a system, method, and program for writing data to a storage medium formatted into a plurality of zones. Each zone has a plurality of addressable sectors, and the innermost zones have fewer sectors than outermost zones. A request is received to write a file to the storage medium. A determination is made of a utilization factor for the file and one zone for the file based on the determined utilization factor. The file is then written to the determined zone.

The Examiner was of the opinion that Dimitri (at Fig. 4, Column 8, lines 31-43) teaches the feature of "thereby providing differentiated classes of storage to one or more clients accessing the system" as required in Applicant's Claim 1. Applicant respectfully disagrees.

Although Dimitri refers to a RAID system, nowhere in the cited reference discusses "providing differentiated classes of storage to one or more clients" that results from aggregating logical block names (LBNs) of storage regions having identical level of performance to a selected section of the LBN space. Referring to Fig. 4 and Column 8, lines 18-43, Dimitri discusses a prior art problem of zone formatting. Specifically, zone constant angular velocity (ZCAV) formatted disks in RAID arrays. ZCAV has to do with the predetermined geometry of a disk and not necessarily its measured performance. In RAID arrays, data from a file is striped across multiple disk drives in the array. If the disk drives are zone formatted, a RAID controller stripes the file to different zones in the disks. If this occurs, then the data transfer rates for different stripes may be significantly different. In such case, the performance of the RAID controller is restricted to the performance of the stripe that takes the longest to write, which is the stripe written to the most innermost zone.

To avoid this problem, Dimitri teaches that a zone would be selected using the zone selection techniques based on a combination of file size and utilization history of the file or controller utilization. After selecting the zone, the RAID controller would then write all data stripes to the same zone on the different disks. This would ensure that the data transfer rate across multiple disks is the same, thereby avoiding a situation where the RAID controller performance is limited to the performance of the most innermost zone of all the disks to which data is striped.

In contrast, Applicant's invention is related to providing differentiated classes of storage to one or more clients accessing the system. As discussed above, clients accessing the system can utilize the storage pool with the appropriate performance level needed to carry out the desired class of service (*See,* e.g., Specification page 11, lines 5-8). For example, one client may utilize a RAID 10 service, and another may utilize a RAID 5 service (*See,* e.g., Specification page 9, line 22 to page 10, line 2). In the process of setting up the system, the storage system administrator can determine which of the performance LBN subspaces should be used to support a particular one of the RAID levels. For example, if Region A is the region of the drives that has particularly good random access I/O performance, it will often be appropriate to allocate it to a RAID-10 set since RAID-10 is also characterized by good random access performance, especially random write performance; the characteristics of the two layers thus reinforce each other resulting in a "Pool A" that has excellent random write performance. (*See,* Specification page 10, lines 3-10). Dimitri has nothing to do with providing differentiated classes of storage as a service to one or more clients accessing the system in response to observed performance levels. Rather Dimitri is solving the prior art problem that has to do with ZCAV formatted disks in RAID arrays.

Moreover, the Examiner was of the opinion that Dimitri teaches the feature of "<u>determine[ing]</u> a level of performance for the plurality of storage locations and <u>partition[ing]</u> the plurality of storage locations into a plurality of regions as determined by their different levels of performance" as required in Applicant's Claim 1. Applicant respectfully disagrees. Dimitri does not determine any level of performance on the storage disk. Rather Dimitri initially assumes that performance would increase by moving data based on the disk geometry, e.g., outer zones having a higher data transfer rate in megabytes per second than the inner zones due to the

greater angular density of sectors on the outer zones. Because the outer zones have a relatively greater data transfer rate, overall disk data transfers rates can be improved by writing the files that are more frequently utilized and larger to the outer zones. Smaller files and files utilized less frequently can be written to a relatively inner zone because slower data transfer rates for such files will not have as adverse an effect on performance as a slow data transfer rate for a larger or more frequently accessed file. For instance, database files or files used by the operating system, such as the file allocation table (FAT), should be written to the outermost zones as they are frequently accessed. Further database files such as a Microsoft Excel file or a Lotus 1-2-3 file can also be quite large and frequently accessed. (*See,* Dimitri Column 6, lines 7-25; and Fig. 4).

In contrast, Applicant's invention recognizes that not all storage locations on a disk have the same performance. For this reason, it is proposed to determine the level of performance, such as access time, reliability of a portion of a disk, and to divide the range of logical block names of the drive into sections rather than simply further using ZCAV information.

Therefore, Dimitri does not anticipate Applicant's invention; as Dimitri also does not disclose a storage device having a plurality of storage locations and a logical block name space for organizing logical block names of the storage locations; or a mapping process configured to aggregate the logical block names of the storage locations in the partitioned regions having an identical level of performance to a selected section of the logical block name space as recited in Claim 1.

Perhaps a practical example would further emphasize the differences between the approach of Dimitri and the Applicants – and the advantages of the latter. Consider a storage system using two different disk drives, each having relatively different performance. A first drive has a relatively faster access time than a second drive. Dimitri's zone-based approach would always group the inner zones of each drive in a RAID array, and group their outer zones together, based on their physical location. However, with Applicants' approach, performance would be first measured, and the inner regions of the slower disk would be grouped with the outer regions of the faster disk.

For these various reasons, the rejection of Claim 1 in view of Dimitri should be withdrawn.

Claims 2-4, 8-10, and 22-27 depend on Claim 1 and therefore are also patentably distinguishable for at least the same reasons described above.

Base Claims 11 and 19 include similar limitations as Claim 1. Therefore, for at least the same reasons described above, Claims 11 and 19 are also patentable.

Claims 12-14, 18, and 28-31 dependent upon Claim 11 and, therefore, are also patentably distinguishable for at least the same reasons described above.

Claim 20 is dependent upon Claim 19 and, therefore, is also patentably distinguishable for at least the same reasons described above.

Accordingly, withdrawal of these rejections is respectfully requested.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By ____________________
David J. Thibodeau, Jr.
Registration No. 31,671
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133

Date: 5/16/08

EXHIBIT A
16/26

TITLE OF INVENTION:

EQLC-PXX-005

Differentiated Storage Pools over a common set of storage devices with multiple RAID levels on each device.

TECHNOLOGY TO PROTECT:

Every patent application is required to end with at least one concise statement, called a claim, of the subject matter that the patent is to protect. To this end, please provide a concise statement of the technology (system or process) you want this patent to protect. Focus on what you want to get a patent on for the company, and do not be concerned that you are not sure it is patentable. We will check the prior art to determine patentability before preparing the application. We will use your statement regarding what you want to protect to determine whether there is an available patent position for your company.

In a block storage environment with multiple devices, differentiated pools of storage are created by applying multiple RAID techniques simultaneously on each device. These pools of storage differ in regards to performance and reliability. Adaptive load balancing can be used to place data in pools to optimize service for a single or multiple consumers of storage services.

For example:

a) Storage devices offer different performance across their LBN space: some LBN ranges have higher bandwidth for read and/or write, some LBN ranges have higher throughput for small and/or random operations, some LBN ranges have slower bandwidth, and some LBN ranges have slower throughput.
b) RAID levels have different performance characteristics for bandwidth and throughput for read and write operations, as well as in both normal mode and degraded operating modes.
c) Device performance differences across LBN space, and performance and reliability differences of RAID techniques, can be combined of to create storage pools that offer different classes of service across a common set of storage devices to the consumers of storage services.

Adaptive load balancing provides the ability to adjust data placement in pools for continuous optimal performance.

The technology to be protected is – minimally – :

1) The use of multiple raid levels concurrently on each device to create differentiated classes of storage service.
2) The use of multiple raid levels concurrently on each device in combination with device LBN performance differences to create differentiated classes of storage service.
3) Load balancing storage services across multiple differentiated classes of storage service:
    a. Automatic, adaptive load balancing of a single or multiple storage services (volumes or file systems) across pools via
        i. Volume based allocation
        ii. Extent based allocation
        iii. Page-based allocation
    b. Manual load balancing of whole volumes

DESCRIPTION OF INVENTION:

Please provide a brief description of your invention, making sure that your description is directed to the technology you want to protect as described in the section above. You can use any format you chose, but experience shows that it is most helpful to provide two or three simple drawings that illustrate the invention. These can be flow charts, functional block diagrams or suitable sketches. Examples of such drawings are provided at the end of this document. We prefer to receive the drawings in electronic form, such as PowerPoint, Visio or SmartDraw. Once the drawings are complete, please add one or two paragraphs that describe <u>what is being shown in the drawings</u>. Please make sure that you identify each element or step shown in the drawings. Consider including a summary of results achieved, features believed to be novel, further experimental work planned, and any additional information.

EXHIBIT A
22/26

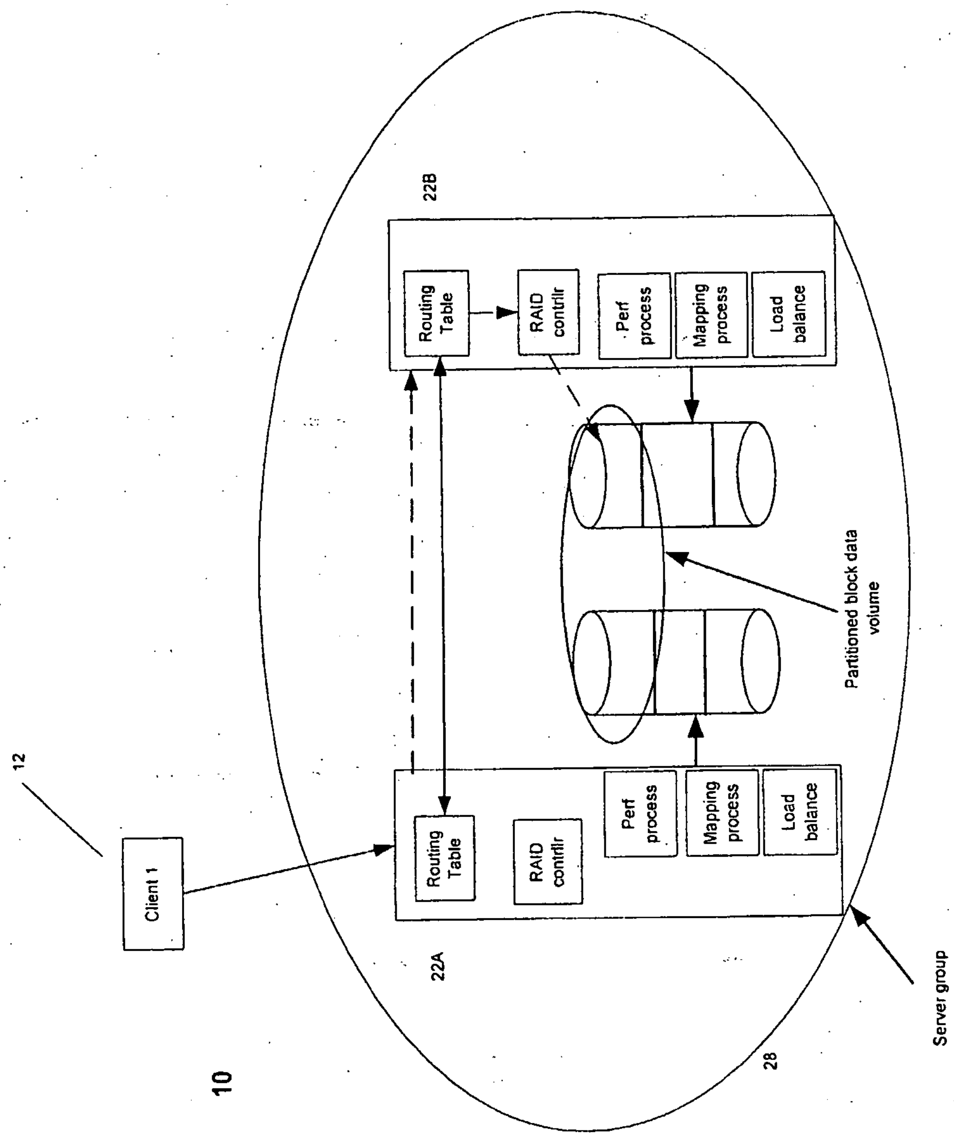

FIG 5

EXHIBIT A
23/26

Diagram 1
Storage device
performance differences
by LBN list

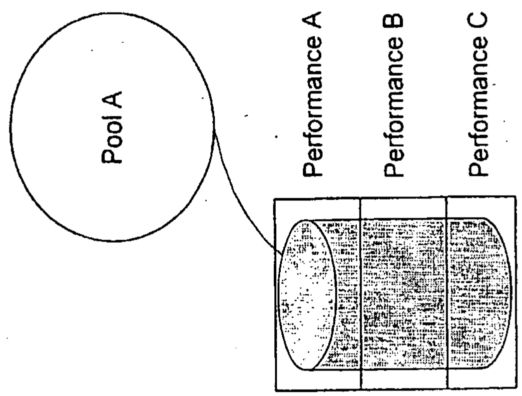

FIG. 6